HP Caliper

A Powerful Itanium® Performance Tool

Curt Wohlgemuth, Steve Williams
HP Caliper Development Team
Hewlett-Packard

Agenda

- Quick overview of Caliper
- Some examples
- New stuff

Quick Overview

- Performance measurement & analysis tool
- Itanium native code
- Linux and HP-UX
- Lots of different measurements using the Itanium PMU
- Per-task (follow processes) and per-CPU (system-wide)
- Measures any application

Some example measurements

Basic command line syntax:

caliper <measurement> [options] application ….

-- or --

caliper <measurement> [options] --attach=<pids>

-- or --

caliper <measurement> [options] --scope=system

e.g.:

caliper scgprof -p sweep3d script_file

caliper fprof --scope=system --duration=60

Measurement types

Totals: ecount, cpu

Profiles: alat_miss, branch_prediction, dcache_miss, dtlb_miss, fprof, icache_miss, itlb_miss

Call graph: scgprof

Traces: pmu_trace

What?
(totaled PMU)
<< 1% intrusion

Where?
(sampled PMU)
~ 1% - 3% intrusion

HowTo - One interesting scenario

- Poor man's call stack near an event
  - Use pmu_trace, sample on DATA_EAR_EVENTS, collect BTB
  - No histogramming, just a raw trace:

```
PMU Trace Buffer (1 of 2), Kernel Thread ID 26462, Samples 1 - 1
--------------------------------------------------------------------------------------------------
         ----------------Data Cache Miss-------------------  ---------BTB Samples----------------
Sample   Address:Slot         Data Address          Latency  Address:slot           Type   FE   BE
Number   (module:function)                                   (module:function)            Misp Misp
--------------------------------------------------------------------------------------------------
     1   0x40000000004767e0:0 0x60000000002b6f90          5  0x40000000004354e0:2   Src    No   No
         (caliper::sqlite3VdbeExec)                          (caliper::initPage)
                                                             0x40000000004355d0:0   Tgt    No   No
                                                             (caliper::getAndInitPage)
                                                             0x40000000004355e0:2   Src    No   No
                                                             (caliper::getAndInitPage)
                                                             0x40000000004360f0:0   Tgt    No   No
                                                             (caliper::sqlite3BtreeCursor)
                                                             0x4000000000436220:2   Src    No   No
                                                             (caliper::sqlite3BtreeCursor)
                                                             0x4000000000476750:0   Tgt    No   No
                                                             (caliper::sqlite3VdbeExec)
                                                             0x40000000004767e0:2   Src    No   No
                                                             (caliper::sqlite3VdbeExec)
                                                             0x4000000000433dd0:0   Tgt    No   No
                                                             (caliper::sqlite3BtreeFlags)
```

HowTo - Tuning Strategy

Step 1 - Identify problem

Algorithmic issues

- scgprof
- fprof

CPU bottleneck

- ecount
- cpu

Step 2 - Pin-point location of problem

- fprof / scgprof
- dcache_miss / dtlb_miss
- icache_miss / itlb_miss
- <customized>

Step 3 - Fix, goto step 1

New feature: cpu metrics

- Measure and report groups of CPU metrics
- Implemented using multiplexed metrics sets
  - requires Perfmon2.2 features
- Great as a “first cut” measurement
- Example metrics sets
  - CPI
  - stalls
  - cache miss data
  - overview: most useful metric sets in one run

New feature: system-wide improvements

- Can use any measurement
- Measure all CPUs, for counts and/or samples
- Samples can be attributed to
  - individual processes
  - individual load modules within processes
  - to no processes at all
- Attribution to modules gives details down to assembly code

New feature: GUI

- Eclipse RCP-based client
- Connect to Caliper server, to
  - perform measurement collections
  - browse Caliper databases
  - see measurement data, with easy drill down

Initial Screen: Follow Red Flags

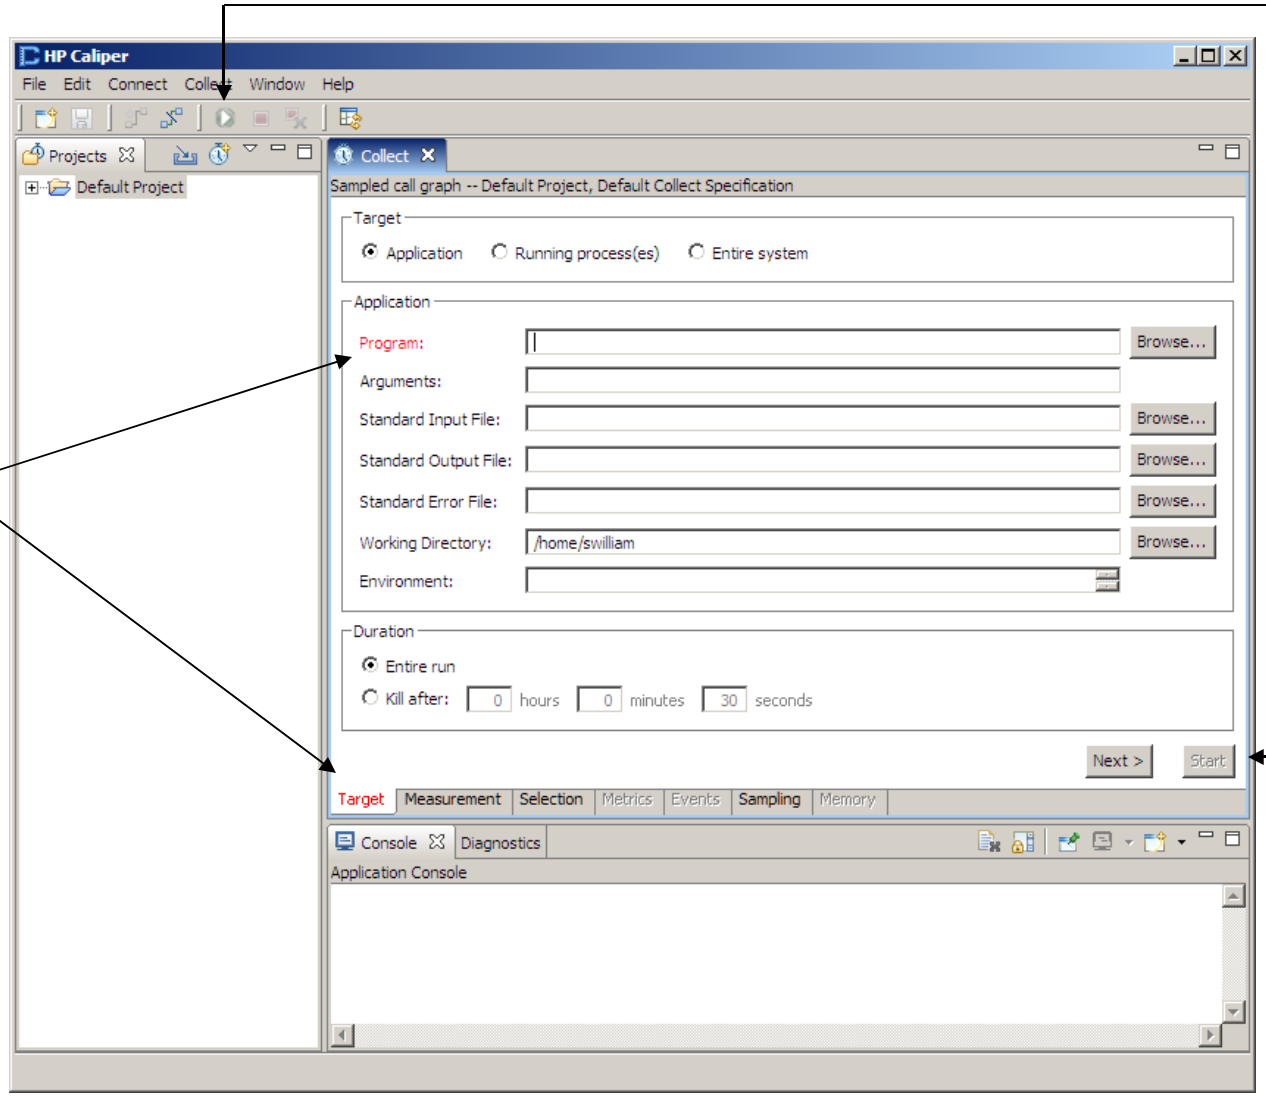

Red Flags

Start Measurement

Choosing Measurements

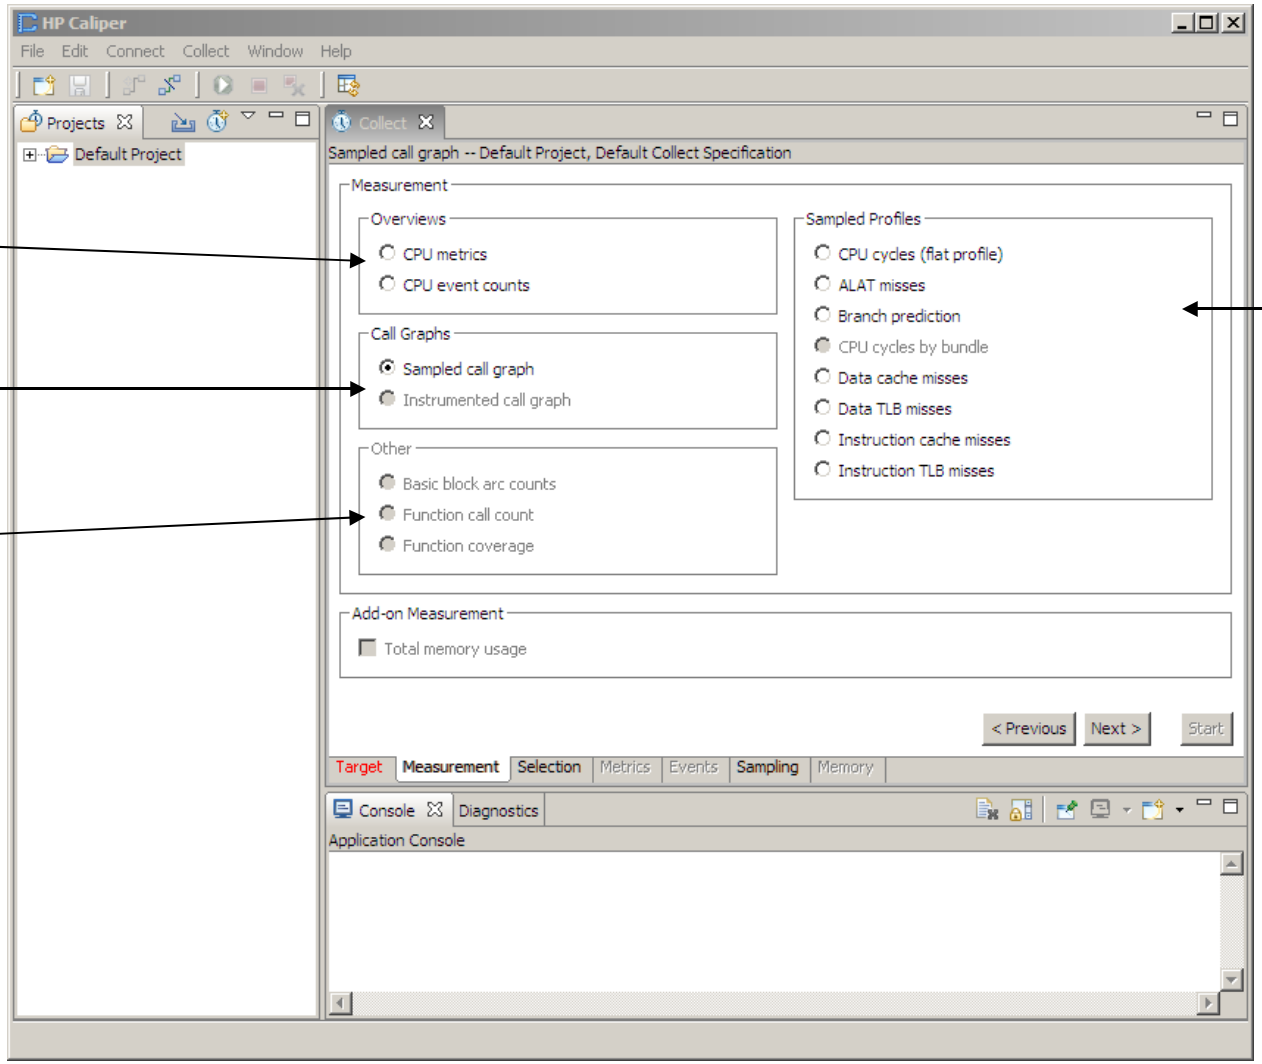

CPU Behavior

Call Graph

Not Available

Profiles

CPU Metrics (L1 Data Cache)

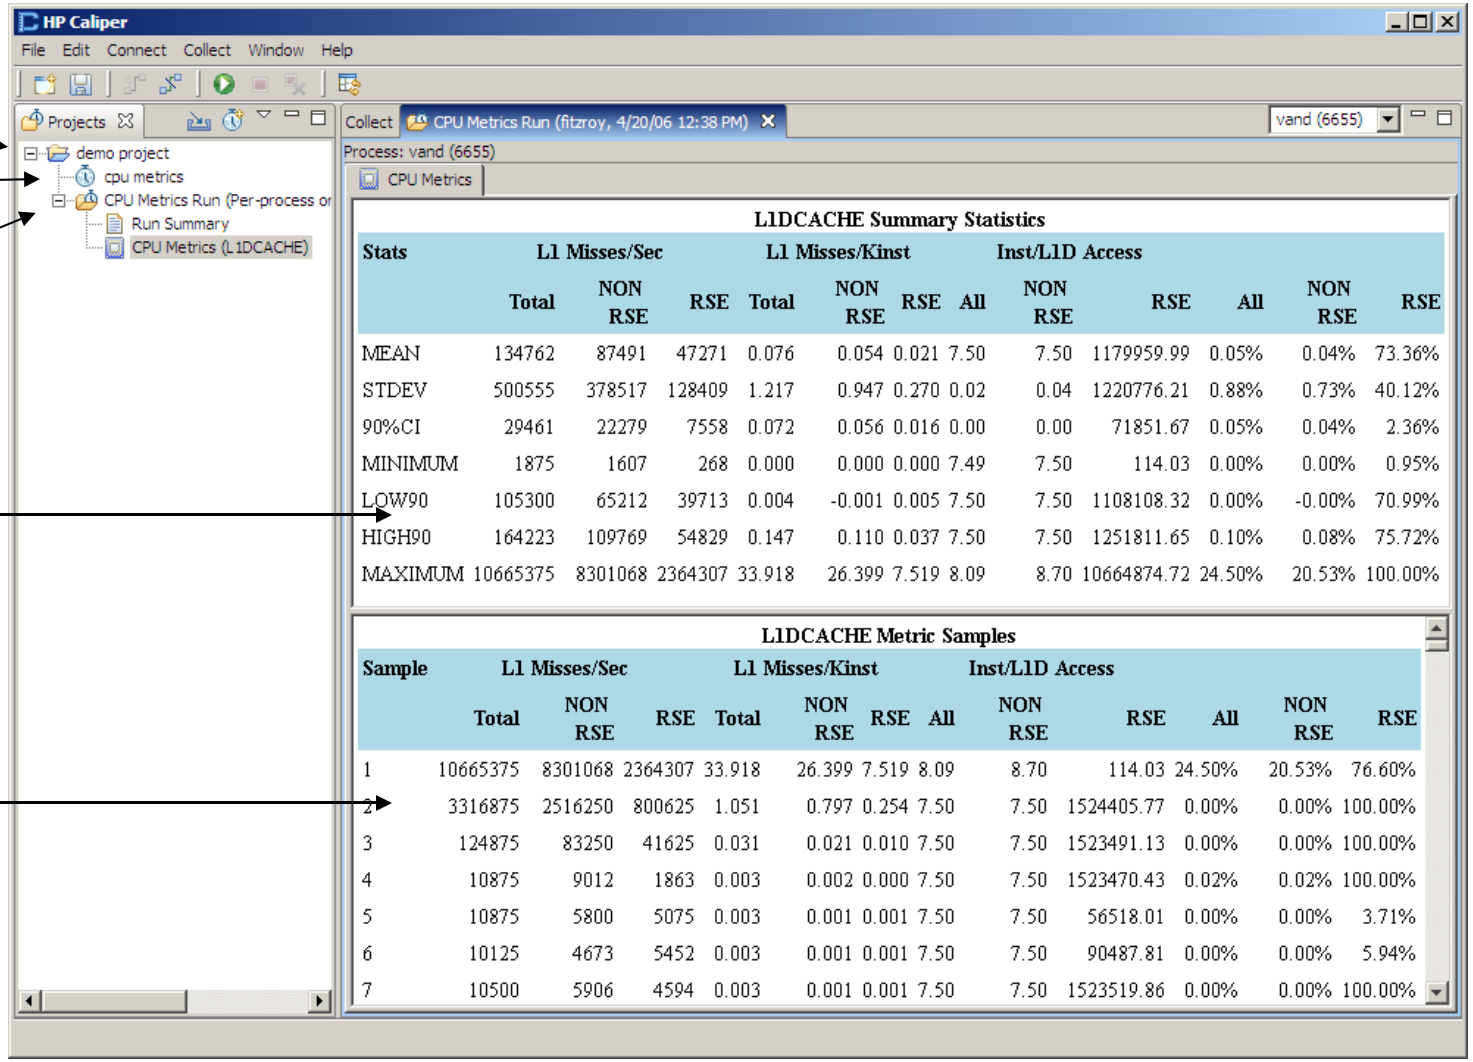

Project

Collection Setup

Collected Data

Summary Statistics

Samples

L1DCACHE Summary Statistics

Stats	L1 Misses/Sec			L1 Misses/Kinst			Inst/L1D Access					
	Total	NON RSE	RSE	Total	NON RSE	RSE	All	NON RSE	RSE	All	NON RSE	RSE
MEAN	134762	87491	47271	0.076	0.054	0.021	7.50	7.50	1179959.99	0.05%	0.04%	73.36%
STDEV	500555	378517	128409	1.217	0.947	0.270	0.02	0.04	1220776.21	0.88%	0.73%	40.12%
90%CI	29461	22279	7558	0.072	0.056	0.016	0.00	0.00	71851.67	0.05%	0.04%	2.36%
MINIMUM	1875	1607	268	0.000	0.000	0.000	7.49	7.50	114.03	0.00%	0.00%	0.95%
LOW90	105300	65212	39713	0.004	-0.001	0.005	7.50	7.50	1108108.32	0.00%	-0.00%	70.99%
HIGH90	164223	109769	54829	0.147	0.110	0.037	7.50	7.50	1251811.65	0.10%	0.08%	75.72%
MAXIMUM	10665375	8301068	2364307	33.918	26.399	7.519	8.09	8.70	10664874.72	24.50%	20.53%	100.00%

L1DCACHE Metric Samples

Sample	L1 Misses/Sec			L1 Misses/Kinst			Inst/L1D Access					
	Total	NON RSE	RSE	Total	NON RSE	RSE	All	NON RSE	RSE	All	NON RSE	RSE
1	10665375	8301068	2364307	33.918	26.399	7.519	8.09	8.70	114.03	24.50%	20.53%	76.60%
2	3316875	2516250	800625	1.051	0.797	0.254	7.50	7.50	1524405.77	0.00%	0.00%	100.00%
3	124875	83250	41625	0.031	0.021	0.010	7.50	7.50	1523491.13	0.00%	0.00%	100.00%
4	10875	9012	1863	0.003	0.002	0.000	7.50	7.50	1523470.43	0.02%	0.02%	100.00%
5	10875	5800	5075	0.003	0.001	0.001	7.50	7.50	56518.01	0.00%	0.00%	3.71%
6	10125	4673	5452	0.003	0.001	0.001	7.50	7.50	90487.81	0.00%	0.00%	5.94%
7	10500	5906	4594	0.003	0.001	0.001	7.50	7.50	1523519.86	0.00%	0.00%	100.00%

Histogram Viewer (Top Level)

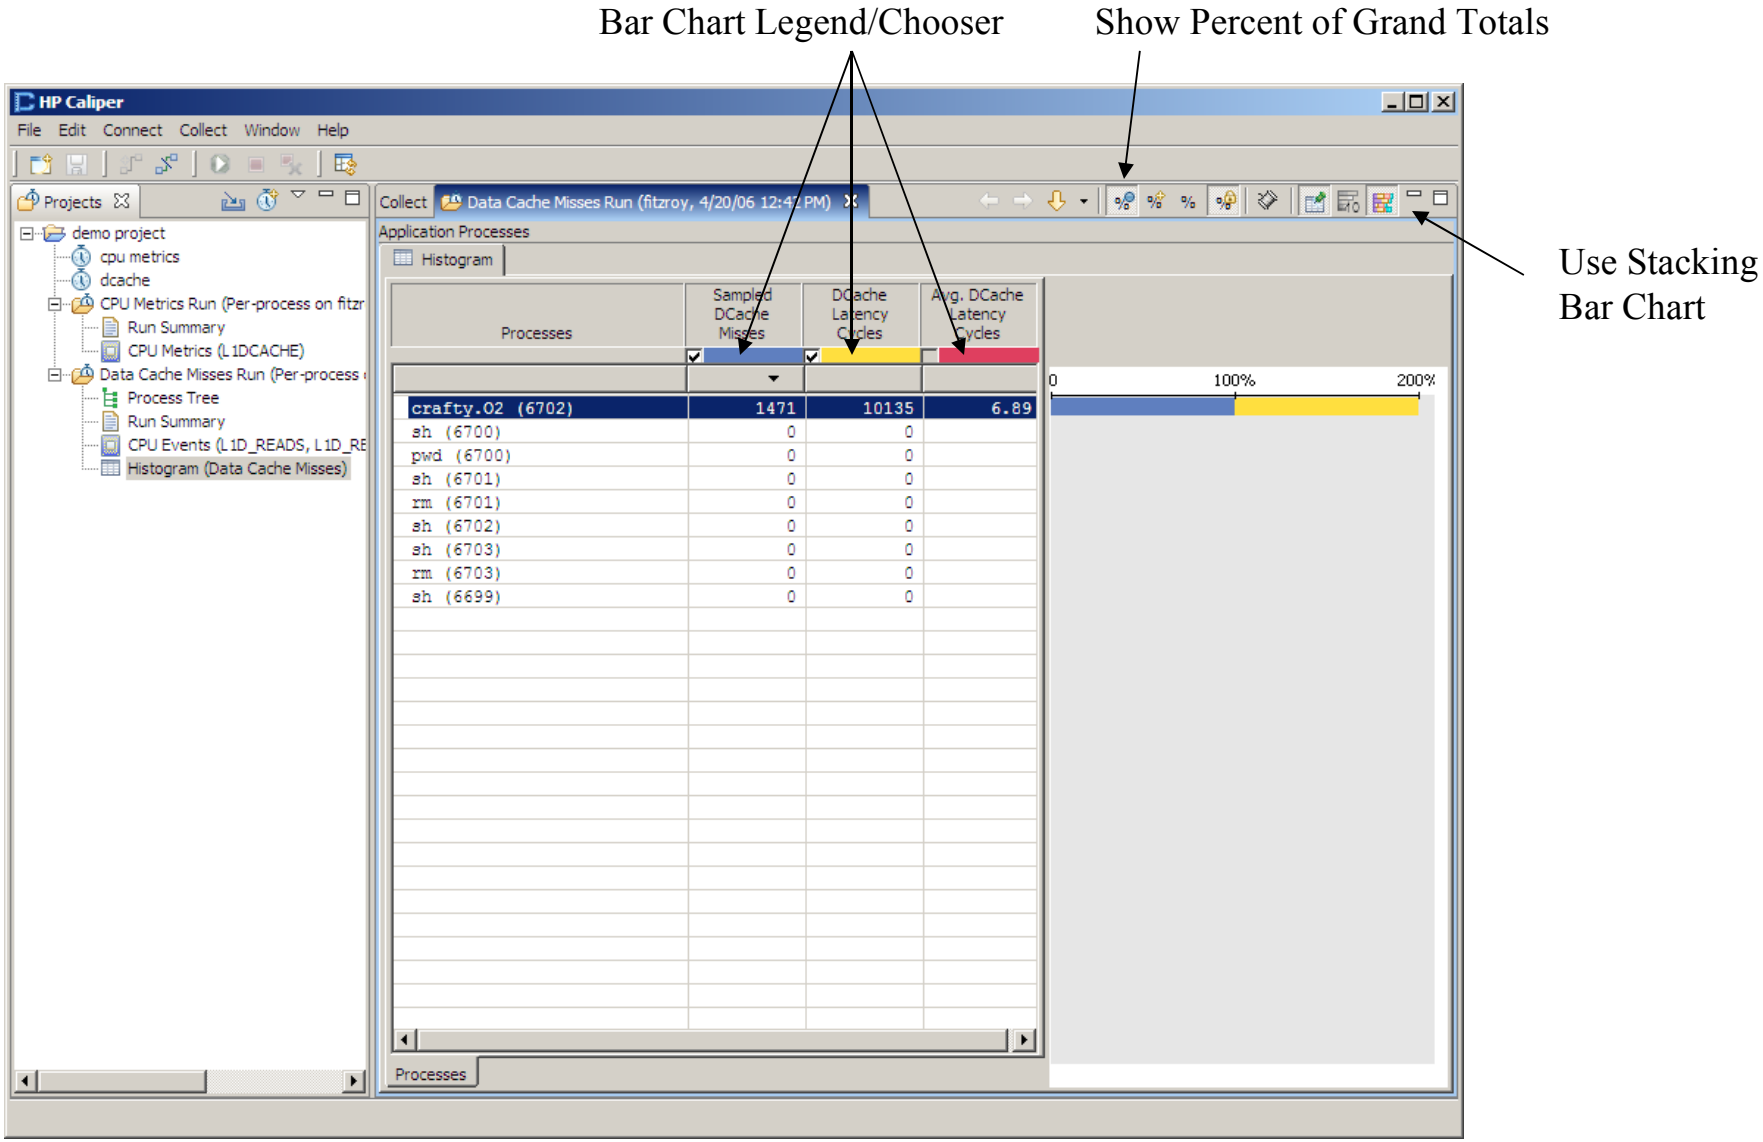

Function Histogram (Data Cache)

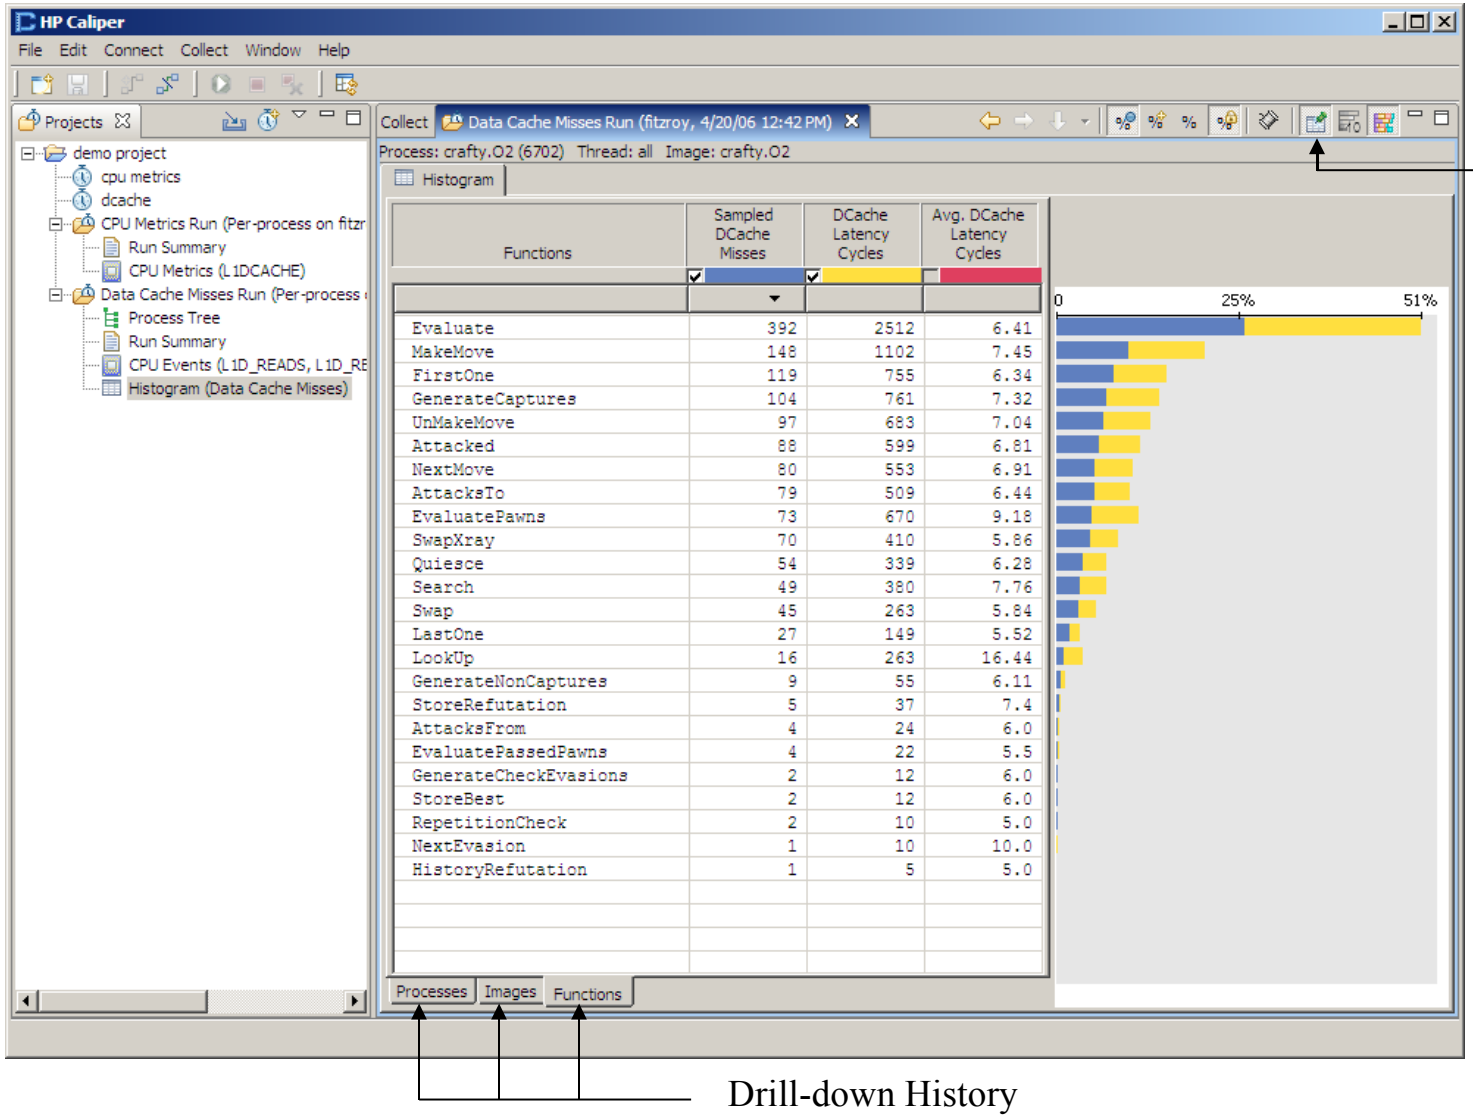

"Pin" Graph to Table

Drill-down History

Disassembly with Overview Graph

Modes: Source, Disassembly with Source, Disassembly

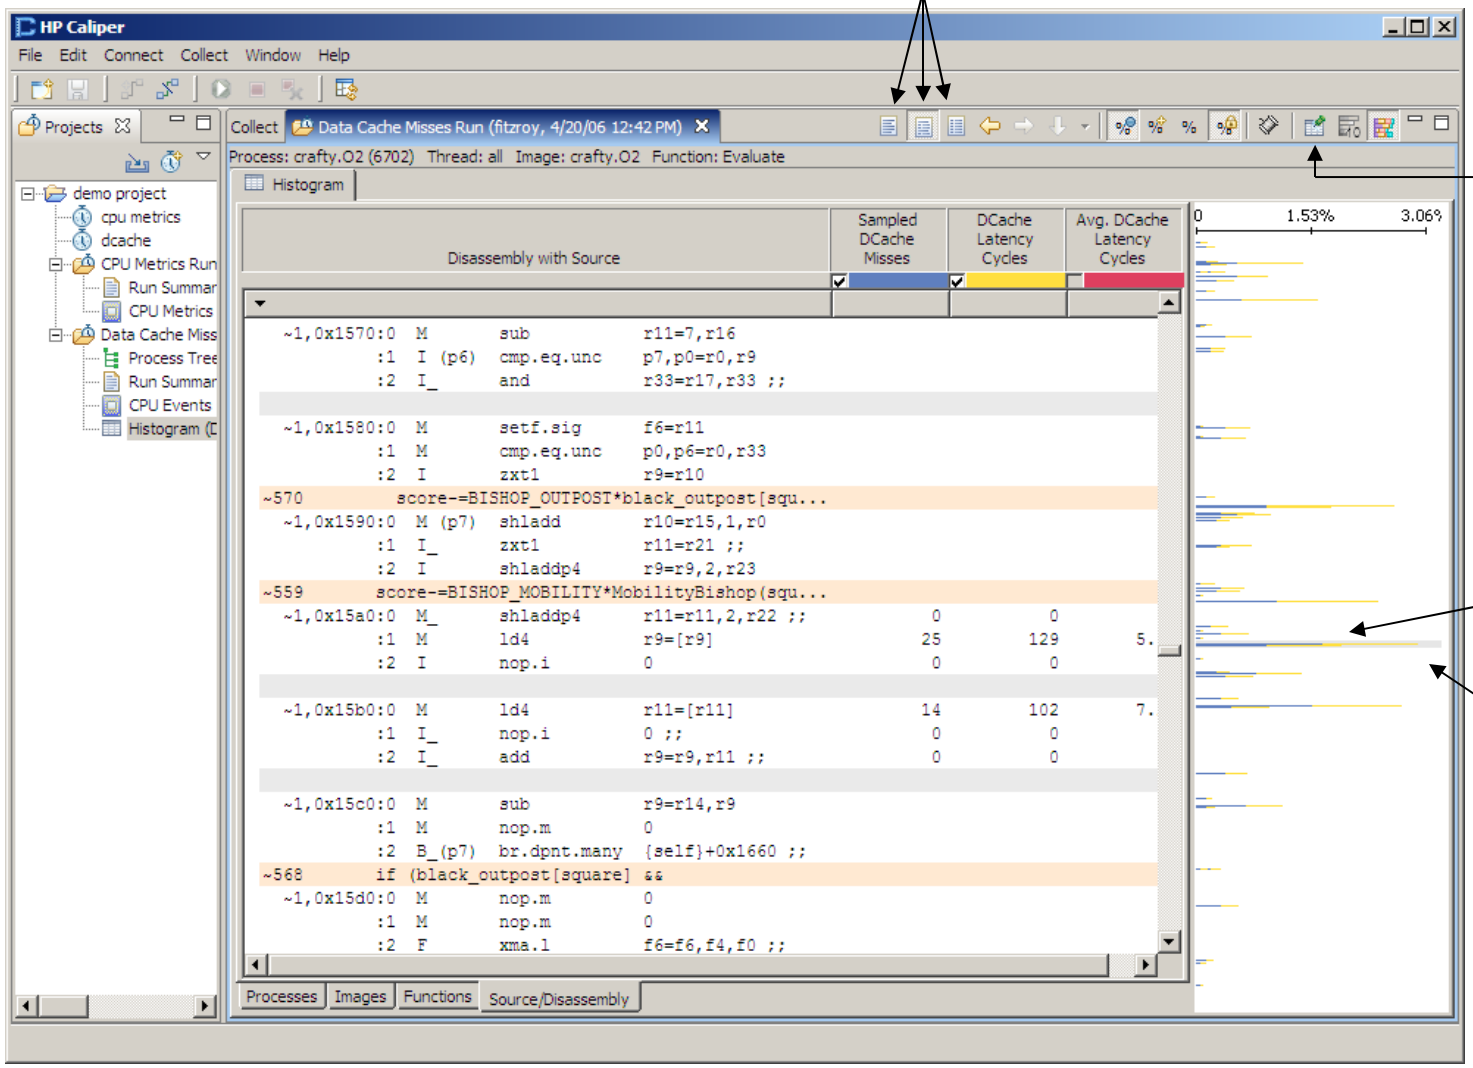

Graph Not Pinned to Show Entire Function

Click Graph to Show Hot Spots in Table

Highlight Shows Table's Current View

Sort by Cache Misses, Pinned Graph

Instruction Bundles Sorted by Data Cache Misses

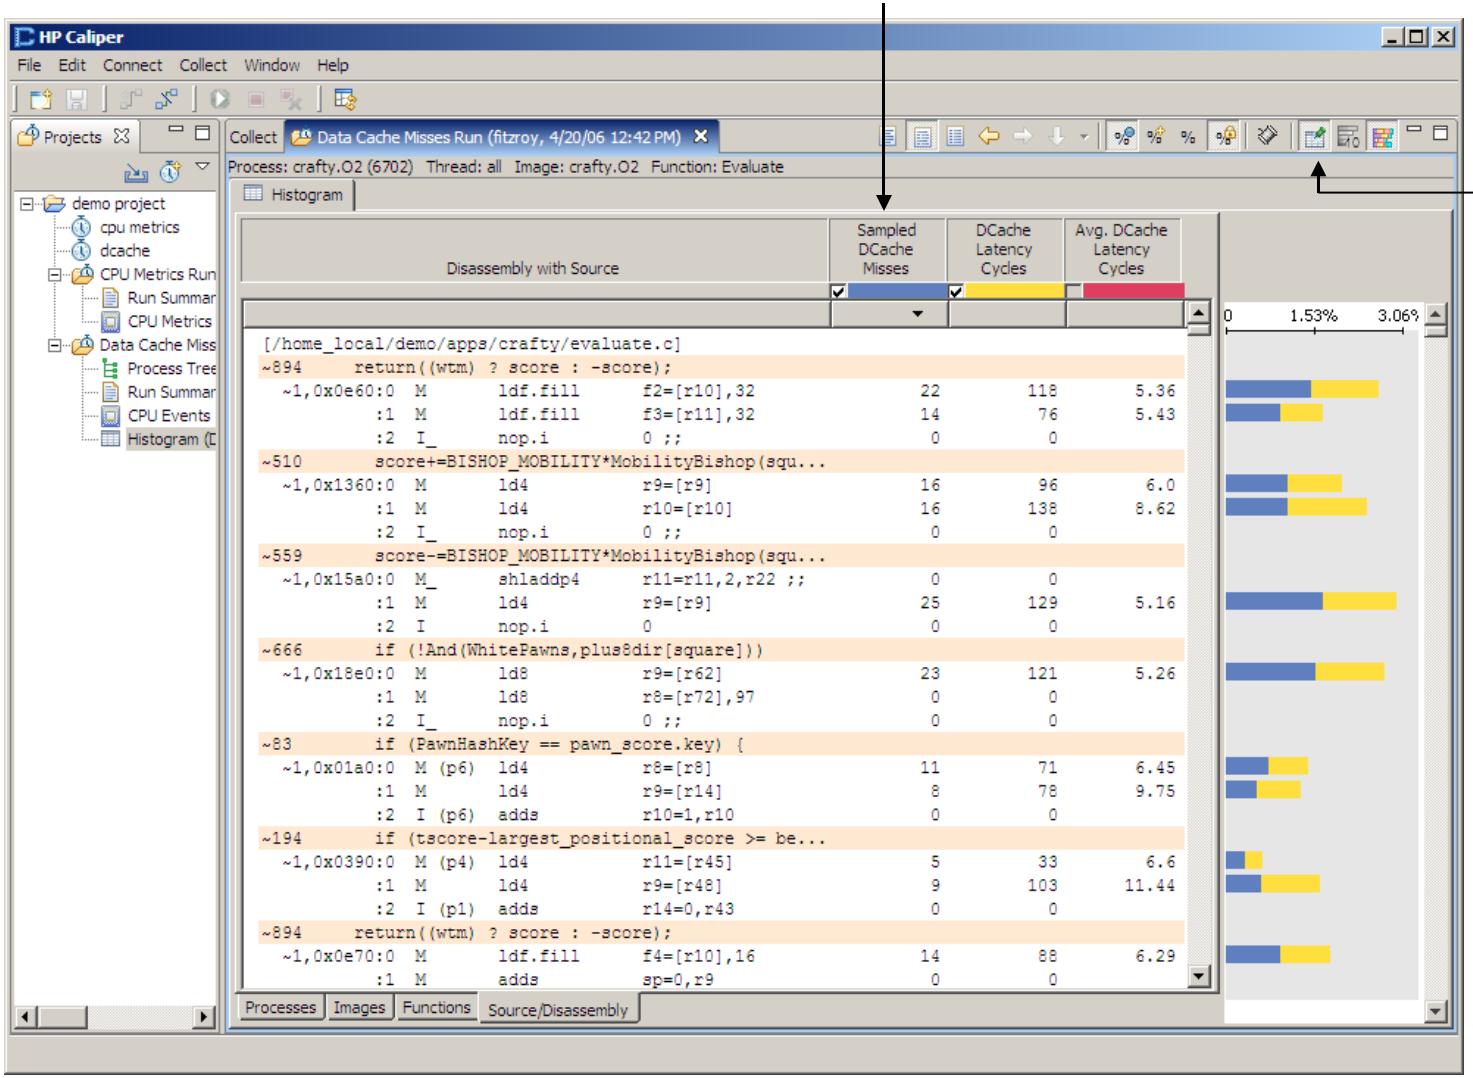

Graph Pinned to Table

Viewing a Process Tree

vfork() Symbol

exec() Symbol

Choose Process to View Its Data

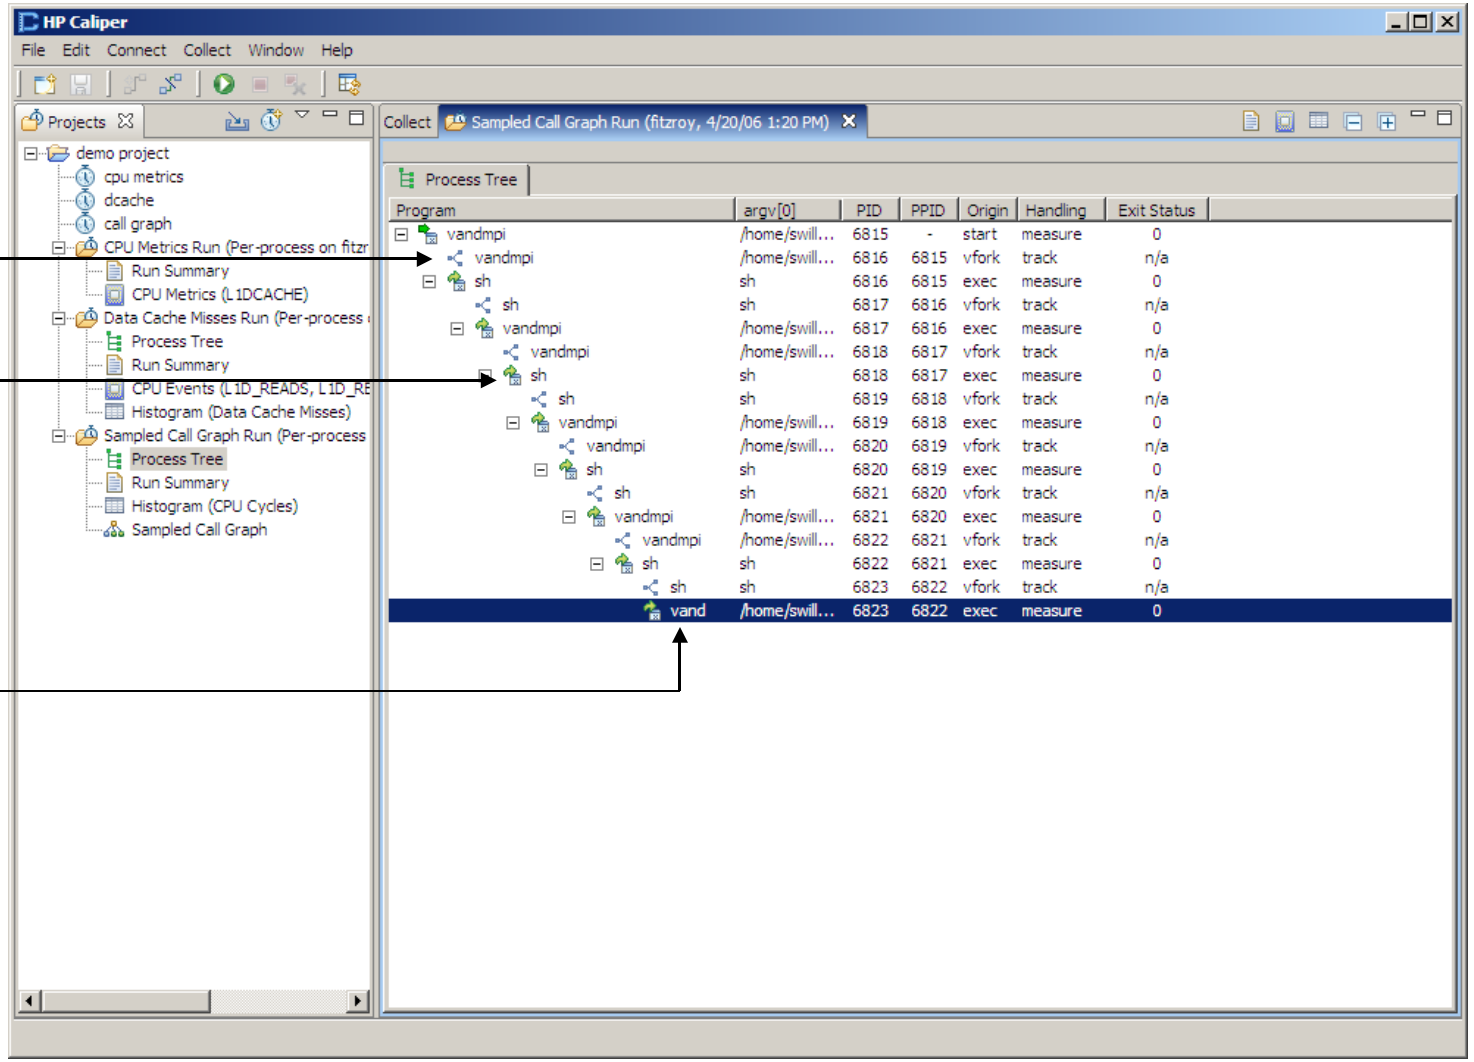

Call Graph Viewer

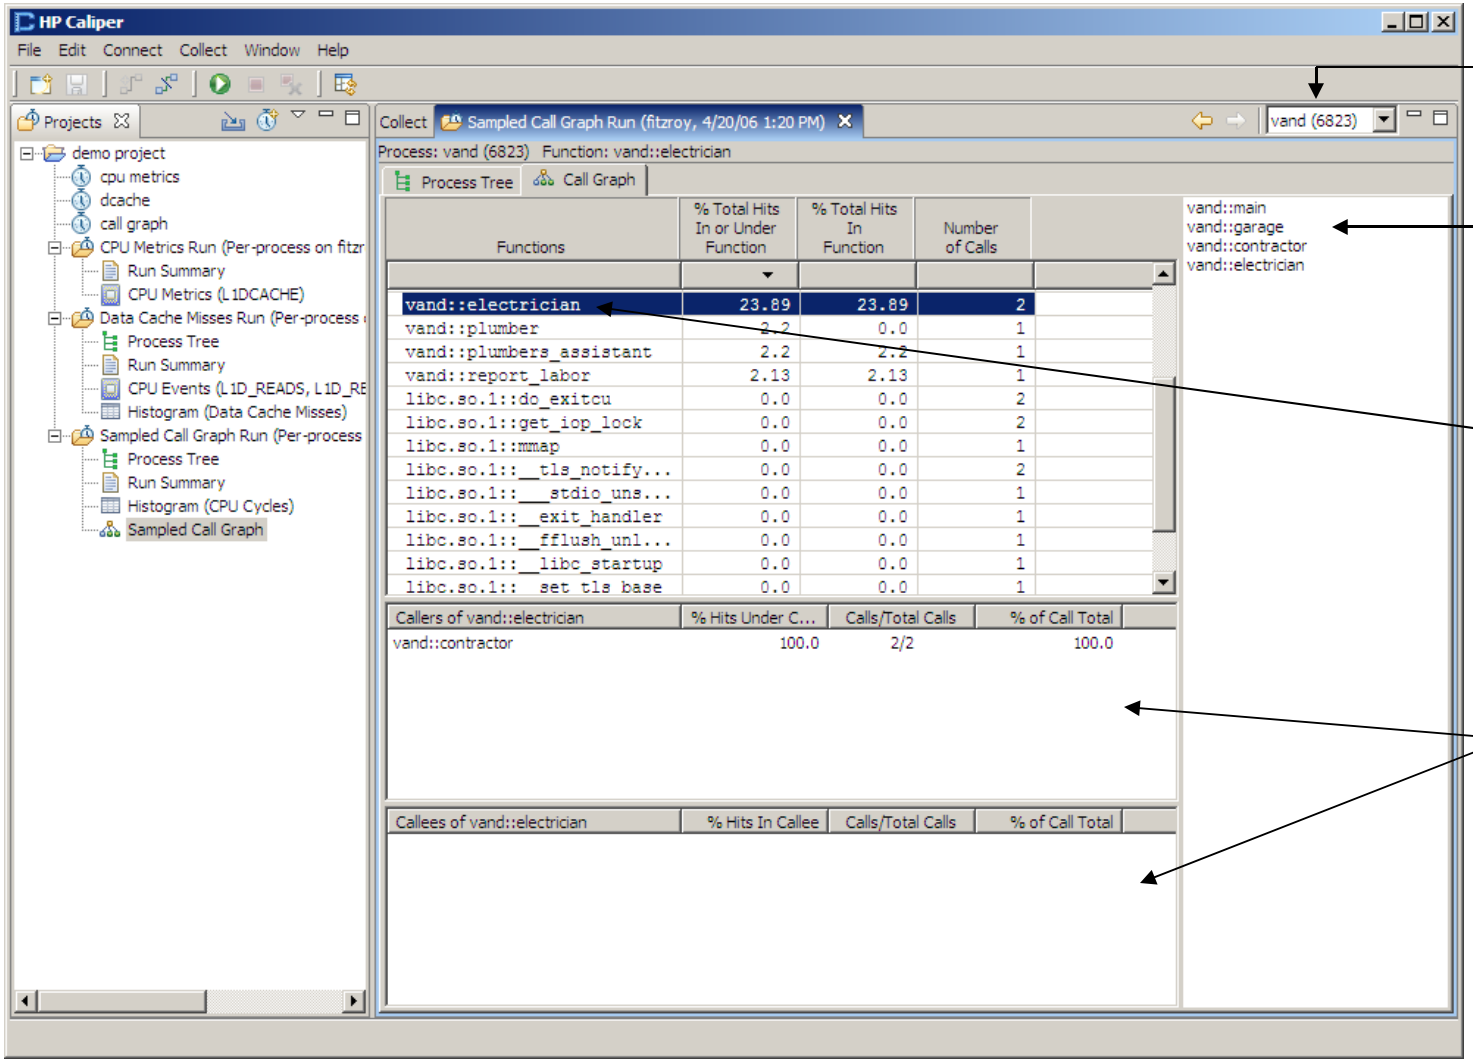

Process Selection

Drill-down History

Selected Function

Callers and Callees

Thanks - Those without whom…

- Processor designers
  - Jim Callister of Intel
- Perfmon subsystem maintainers
  - Stephane Eranian (Linux)
  - Dan Truong (HP-UX)

Resources

- caliper-help@cup.hp.com
  - Caliper “help” mailing address
  - Caliper questions, bug reports, enhancement requests
- http://hp.com/go/caliper
  - Caliper external website
  - Downloads, documentation, tech notes
- http://devresource.hp.com
  - Developers’ Resource website
  - Training webinars
- Gelato meeting, May 2005 Caliper slides
  - http://www.gelato.org/pdf/may2005/gelato_may2005_caliper_gouriou_hp.pdf

Questions